

## Fairness in Algorithmic Policing

**Abstract:** Predictive policing, the practice of using of algorithmic systems to forecast crime, is heralded by police departments as the new frontier of crime analysis. At the same time, it is opposed by civil rights groups, academics, and media outlets for being “biased” and therefore discriminatory against communities of color. This paper argues that the prevailing focus on racial bias has overshadowed two normative factors that are essential to a full assessment of the moral permissibility of predictive policing: fairness in the social distribution of the benefits and burdens of policing as well as the distinctive role of consent in determining fair distributions. When these normative factors are given their due attention, several requirements emerge for the fair implementation of predictive policing. Among these requirements are that police departments inform and solicit buy-in from affected communities about strategic decision-making and that departments favor non-enforcement-oriented interventions.

**Key words:** Fairness, Consent, Policing, Algorithms, Artificial Intelligence, Data

## **Introduction: a new frontier of policing**

In 2011, the LAPD began an experiment in crime fighting when they implemented a computer program called “PredPol” to anticipate the timing and location of crime. PredPol is an algorithmic system that takes data about the type, location, and time of crimes as inputs and produces predictions about when and where future crimes will occur. PredPol was used by the LAPD to forecast future property crime—vehicular theft and theft from a vehicle—at highly specific locations and times during a police officer’s patrol shift. PredPol forecasts are displayed on a computer screen as “heat maps” consisting of 500x500 square foot boxes, each indicating high-crime-risk areas in an officer’s beat. Additional data from recent crime reports can be incorporated each day to update the algorithm’s forecast. While the LAPD discontinued its PredPol program in the spring of 2020, PredPol has become one of the most widely used pieces of predictive policing software in the United States (Miller 2020). Dozens of police departments across the country have used or are currently using it to allocate patrols (Jouvenal 2016). The lure of predictive policing is that a data-driven approach might allow police departments to uncover crime patterns unnoticed by human analysts while avoiding human bias in police resource allocation (Selbst 2017). A further hope underpinning the adoption of PredPol is that mere visibility in high-crime areas at specific locations and times can help police to deter crime without making direct contact with offenders (Brantingham, Valasik, and Mohler 2018).

Predictive policing has received withering criticism in the popular press (Haskins 2019), by academics (Shapiro 2017; Lum and Isaac 2016; Richardson, Schultz, and Crawford 2019), and by community organizations like the ACLU and NAACP. One widespread

criticism is that predictive policing systems make their predictions on the basis of data that is “biased” against racial minorities. This bias leads algorithmic crime forecasts to recommend inundating minority communities with police, when in reality those resources would be more efficiently spent elsewhere. In this paper I evaluate this widespread criticism, argue that it is inconclusive, and explore a new way forward in the debate about the fairness of predictive policing. I propose that predictive policing can be unfair even if it is unbiased. The threat of unfairness is especially acute in cases where predictive policing produces unequal burdens among citizens, and those citizens who are disproportionately burdened by predictive policing have legitimately refused consent to the benefits it provides.

Section 2 addresses the dominant criticism of predictive policing: that predictive policing systems make use of racially “biased” data, where data is racially biased insofar as it reflects racist policing practices rather than the rate of crime at particular locations or by particular people. According to this criticism, predictive policing algorithms systematically overrepresent the proportion of criminal activity in predominantly black communities, usually to the detriment of those communities. I argue that this objection is not decisive. In any particular case, the argument that predictive policing systems are racially biased is often inconclusive.

In section 3, I show that the dominant focus on algorithmic bias omits important ethical considerations from the conversation about the permissibility of predictive policing. Specifically, the emphasis on bias implies that a predictive policing program is unfair only if it makes use of biased data. I illustrate one way in which a *non*-racially biased predictive policing program might be unfair: the program might impose an unequal burden of law

enforcement on innocent members of some racial groups for the sake of greater security gains on aggregate. This policy could be both unbiased—in the relevant sense—and effective while also being unfair. Call this the “Unequal Burdens Argument.” I then devote significant space to defending the Unequal Burdens Argument from an objection. According to the objection, disproportionately burdening minority communities in this way is not unfair as long as those burdens are outweighed by disproportionate benefits in the form of crime reduction (Boonin 2011; Risse and Zeckhauser 2004). In responding to this objection, I argue that individuals possess the normative authority to exclude their well-being from being among the considerations that justify risky interventions. Because predictive policing is a risky intervention, the benefits of predictive policing for affected communities count morally in favor of the practice only if members of the affected communities endorse or consent to policing by algorithm. The upshot of this discussion is that achieving fair predictive policing requires going beyond “de-biasing” algorithms. Predictive policing must be welcomed by targeted communities to be fair. Section 4 concludes by offering modest practical guidance for the implementation of fair predictive policing programs.

## **1 Preliminaries**

### *What is predictive policing?*

While there is no uniform definition of predictive policing, I will follow a definition offered by Albert Meijer and Martijn Wessels: “Predictive policing is the collection and analysis of data about previous crimes for identification and statistical prediction of individuals or geospatial areas with an increased probability of criminal activity to help

developing policing intervention and prevention strategies and tactics” (Meijer and Wessels 2019). This definition includes both *persons* and *places* as potential targets of prediction by predictive policing systems. Some predictive policing systems, such as Chicago Police Department’s “heat list” and the Intelligence-led Policing Program used by the Pasco County Sheriff’s Office in Florida (“ILP Manual” 2016) identify individuals who are at a high risk of being a victim or offender (Perry et al. 2013; Kump et al. 2016; Ferguson 2017). Systems like PredPol and Hunchlab, on the other hand, only make “place-based” predictions about future crime. These predictions inform police departments seeking to efficiently allocate police resources. If a place or person is deemed high risk, then the police will pay the place or person greater attention.

Before predictive policing systems, crime analysts would generate their own “heat maps” of criminal activity, based on expert knowledge, intuition, and experience in the field. Predictive policing systems enjoy two advantages compared with this old method: the systems can identify patterns in crime data that would go unnoticed by human crime analysts, and the systems are less susceptible to biased “hunches” about where crime is likely to occur. These systems also provide fine-grained predictions, which allows departments to target different types of crime in different places and times throughout a patrol officer’s shift (Shapiro 2017: 459). Predictive policing therefore enables departments to deploy resources “more accurately in place and time” (Meijer and Wessels 2019: 1033).

*Excessive use of force*

In order to isolate the moral problems with predictive policing, I am going to bracket some issues related to systemic police harassment and abuse of people of color. This is not to suggest that these issues are not important—they are—or that they do not bear on the moral permissibility of predictive policing—they do. If police abuse members of the communities to which they are deployed, especially communities of color, this constitutes a powerful moral reason not to deploy them there. But my interest in this paper is to understand whether there is something morally problematic about using predictive policing to conduct policework that is otherwise permitted. For example, many critics of predictive policing agree that police are permitted to conduct vehicle and foot patrols of neighborhoods, so long as officers on patrol do not harass or abuse residents. What critics object to is the practice of *deciding where* to deploy those patrols on the basis of predictive policing forecasts.

### *Police legitimacy*

“Police abolitionism,” as it’s sometimes called, is gaining popularity in the United States. Some of the most outspoken critics of predictive policing such as the Stop LAPD Spying Coalition are police abolitionists (Moravec 2019). For various reasons, police abolitionists oppose the use of a state police force to enforce the laws. In this paper I will not address general objections to the institution of policing as the dominant form of law enforcement of the sort found in Alex Vitale’s recent book *The End of Policing* (Vitale 2017). For the sake of argument, I will set police abolitionism aside and assume that legitimate states are justified in using a domestic police force to enforce the law, at least in cases where those laws are worthy of being enforced. I further assume that one legitimate aim of policing

is the apprehension of individuals who violate the law and the deterrence of individuals from violating the law in the first place, when those laws are worthy of being enforced. By assuming that police abolitionism is false I am not assuming that particular policing practices are beyond reproach. Even if we grant that the institution of policing is legitimate, we may yet question whether a particular policing practice is morally justified.

### *Racial profiling in disguise*

Finally, I will grant that police departments use programs like PredPol in a good faith effort allocate resources to areas with the greatest risk of crime. In granting this, I also grant that predictive policing is not an attempt by police departments to disguise the targeting of racial minorities. In granting these assumptions, I aim to be as charitable as possible to advocates of predictive policing. If I can successfully argue that predictive policing can be unfair, and hence impermissible, even when the above assumptions are granted, I will have provided an argument that is harder to defeat than one that depends on rejecting some of them.

## **2 Is predictive policing effective? Is it biased against racial minorities?**

In this section I discuss evidence for the efficacy of predictive policing at reducing crime, and then I assess the empirical case for the claim that predictive policing is biased against racial minorities. When it comes to evaluating the efficacy of a new policing technology or practice, what matters is whether subject of evaluation is *more effective than* the status quo, as well as other alternatives to the status quo that we might adopt instead, at preventing crime. I will therefore understand a policing policy or strategy to be effective,

other things being equal, if it is better than feasible alternative policies at achieving the legitimate aims of policing. Because I am assuming that the legitimate aims of policing are the apprehension and deterrence of offenders who violate laws that are worthy of being enforced, a policing practice or technology is *effective* if it is better than feasible alternatives at deterring and apprehending individuals who (would otherwise) violate laws that are worthy of being enforced.

*Why predictive policing seems effective in principle*

A simple observation made by criminologists in the late 1980s, and since confirmed by a number of studies, supports the claim that predictive policing is, in principle, an effective crime prevention strategy: crime is connected to place (Pierce et al. 1988; Sherman, Gartin, and Buerger 1989; Weisburd et al. 1993). Crime is concentrated not only within neighborhoods but within specific blocks or groups of blocks within neighborhoods. A non-algorithmic predecessor of predictive policing called hot spots policing has long acknowledged the connection between crime and place. While hot spots policing is itself controversial, the effectiveness of hot spots policing is supported by significant evidence (Weisburd and Eck 2004; National Research Council 2004; A. Braga, Papachristos, and Hureau 2012; A. A. Braga, Papachristos, and Hureau 2014).

There exists little empirical literature to date on predictive policing's effectiveness (Boba Santos 2019; Meijer and Wessels 2019). Still, policing by algorithm has enjoyed some successes. A recent data-driven policing project in Atlantic City found a significant decrease in crime rates. The project employed "Risk Terrain Modelling" (RTM), "a method of spatial risk analysis used to assess spatial patterns of crime and diagnose how features of a

landscape interact and overlap to create unique crime settings” (Caplan, Kennedy, and Drawve 2017: 1). Using these forecasts to generate “risk terrain maps,” the Atlantic City PD targeted high risk places with extra police presence. Local civic organizations and city government also participated by targeting high-risk areas with security improvements. Even though these priority areas constituted only 1% of Atlantic City’s overall land area, researchers found that targeting these areas correlated with 36% decline in crime in Atlantic City in 2017 compared with 2016. A study of the New York City Police Department’s predictive policing software concluded that the predictive algorithm was more than twice as accurate as human analysts at predicting future crime (Levine et al. 2017). So, some recent evidence suggests that some predictive policing algorithms are more accurate than human crime analysts. Other studies have shown no difference in the performance of the predictive algorithm and human analysts (Perry et al. 2013).

*Is predictive policing racially biased?*

In spite of predictive policing’s early promise, a chorus of activists, journalists, and academics have argued that it is ineffective because it is unduly biased against people of color. What we can call the “Self-fulfilling Prophecy Objection” to predictive policing is described in a number of places (O’Neil 2016; Lum and Isaac 2016; Ferguson 2017; Selbst 2017; Ensign et al. 2018; Lau 2020). According to this objection, if officer decisions about where to patrol and whom to investigate are racially biased, we should expect more arrests of members of the targeted racial group. When officer decisions are racially biased in a way that disfavors blacks, we should also therefore expect the crime forecasts of algorithms that are trained on arrest data to reflect this biased behavior of police officers on patrol. These

systems will direct officers to patrol black communities out of proportion to the actual rate of criminal activity in those communities.

Predictive policing leads to a self-fulfilling prophecy or feedback loop when police departments respond to forecasts by sending more police patrols to black communities. Greater police attention will yield more arrests in the targeted location, which will be used to update the algorithm's forecasts, which will confirm the earlier predictions.

The Self-fulfilling Prophecy Objection is one instance of the more general “ratchet effect” of police profiling described by economist Bernard Harcourt.

[I]f the police dedicate more resources to investigating, searching, and arresting members of a higher-offending group, the resulting distribution of arrests (between profiled and nonprofiled persons) will disproportionately represent members of that higher-offending group. The basic intuition is that policing is like sampling: when the police profile frequent offenders, they are essentially sampling more among the members of the higher-offending group (Harcourt 2006: 147).

Even if the initial crime information is accurate, responding to that information rationally by increasing the amount of police attention paid to the area can *ratchet up* future police scrutiny to a point at which that scrutiny is no longer justifiable. In the case of predictive policing, the profiled group consists of black individuals living in areas with large numbers of arrests. The disproportionate sampling consists of officers making disproportionately more arrests in those profiled high crime areas. These arrests generate the data that is fed back into the predictive policing algorithm, leading to a rate of arrests in the

profiled areas that exceeds the actual rate of criminality in those areas. Of course, the ratchet effect is amplified if the initial data overrepresents crime in black communities. In this case, predictive policing would be ineffective *because* it is biased.

The Self-fulfilling Prophecy Objection targets predictive policing's effectiveness, but it is a short step from this objection to the conclusion that predictive policing is also unfair; for all sides of the debate will agree that it is unfair to unequally burden minority communities with the costs of law enforcement, at least when crime would be more effectively prevented if patrols were directed to places where the crime is actually occurring. This is precisely what the Self-fulfilling Prophecy Objection purports to establish.

*Why predictive policing is (at least sometimes) not biased*

The Self-fulfilling Prophecy Objection depends on the claim that the input data of predictive policing systems is prone to bias. Yet it is not clear that this claim can be sustained, because some predictive policing systems do not rely on arrest data to make their forecasts (Mohler et al 2015; Ferguson 2017; Brantingham et al 2018). Jeffrey Brantingham, whose academic work was used to create PredPol, writes of the system,

the majority of hotspot and place-based predictive policing algorithms focus not on arrests, but on crimes predominantly reported to the police by the public (e.g., robbery, burglary, assault). Thus, the goal is to send police resources to areas where crimes have been reported by victims, thus preventing future crimes in those areas. While a feedback loop for reported crime may be possible, in this case the self-reinforcement is toward places

where citizens are placing calls for service (Brantingham, Valasik, and Mohler 2018: 3).

As the quote from Brantingham et al suggests, whereas arrests often reflect only discovered crime or suspicion of crime by officers, calls for service are more likely to reflect an actual crime occurrence that has been witnessed or discovered by a member of the public. Data on emergency calls for service are by no means flawless methods of tracking crime rates, but they do not in any obvious way lead to a self-fulfilling prophecy (Black 1970; Gutierrez and Kirk 2017).

Citizen crime reporting can of course show racial bias (Tashea 2016). However, it is unclear that racial bias affects reporting of the sorts of crimes targeted by predictive policing systems. For example, PredPol was used by LAPD to target motor vehicle theft. Because vehicular theft is typically discovered by the victim after the thief has disappeared with the vehicle, the race of the suspect cannot influence the victim's choice about whether or how to report the crime.

Calls for service from predominantly non-white communities could *underrepresent* the extent of property or violent crime in those places. A recent Cato Institute survey found that while 78% of white Americans say they would “definitely” report a violent crime they witnessed, only 54% of African Americans and 57% of Hispanics are so confident (Ekins 2016). Blacks and Hispanics also report a greater degree of legal cynicism—a lack of confidence in the justice, legitimacy, and responsiveness of agents of the law—than whites. This may affect their disposition to report crimes to the police (Kirk and Papachristos 2011). Residents of underprivileged communities are sometimes hesitant to report crimes because

of perceived police harassment or fear that their identities will be revealed to the perpetrators (Anderson 2000). Therefore, a predictive policing system that makes predictions on the basis of community-generated calls for service might *underestimate* crime in predominantly black communities, leading to less police attention there than is warranted by actual crime rates. The popular Self-fulfilling Prophecy Objection to the effectiveness of predictive policing is therefore inconclusive.

The fundamental limitation of the Self-fulfilling Prophecy Objection to predictive policing is that it hitches the case against predictive policing to the empirical claim that predictive policing systems are trained on racially biased data. While this is a legitimate concern, it is difficult to establish that racially biased data is being used in particular cases.

### 3 Fairness in predictive policing

#### *Unfairness and bias*

If predictive policing turns out to be effective and unbiased, then it is not obvious how it could be unfair. Some critics of predictive policing seem to concede this. Consider Andrew Selbst's articulation of a fairness-based objection to predictive policing.

The potential for harm stemming from racially imbalanced outcomes is the harm resulting from having more police in a neighborhood that is **unfairly** **maligned as having more crime**... Thus, on a systemic level, over-policing nonwhite neighborhoods does present a fairness harm (Selbst 2017: 130).

As indicated by the bolded text, Selbst takes the claim that predictive policing is unfair to turn on the empirical fact that predictive policing overestimates crime in black communities.

Over-policing in black communities is unfair, according to Selbst, precisely because there are too many police there, given the actual amount of crime in those communities. This suggests that predictive policing is *not* unfair, in the sense described by Selbst, as long as forecasts are accurate reflections of the amount of crime in black communities. I will now argue that mere absence of bias does not entail fairness. Predictive policing can be unfair even when it is unbiased, accurate, and based on sound empirical and theoretical rationales.

### *The Unequal Burdens Argument*

To begin to understand why predictive policing can be unfair, even when it is unbiased and effective, it is helpful to consider another policing practice—racial profiling—which is widely criticized as being unfair to black Americans. A comparison between predictive policing and racial profiling can be illuminating for several reasons. First, predictive policing and racial profiling are both methods of allocating police resources and, therefore, police attention. Second, racial profiling has been criticized in the media and academic scholarship as being unfair to racial minorities. Third, some fairness-based objections to racial profiling do not turn on the question of whether racial profiling is effective at reducing crime.

The simplest defense of racial profiling is that crime statistics show that some racial groups commit some crimes at a disproportionately high rate, given their representation in the general population. If it is rational for police to deter those crimes, and to apprehend individuals who commit them, then it is rational for police to disproportionately investigate members of those racial groups (Boonin 2011: 308–26).

Opponents contend that the rationality of a practice does not equate to moral permissibility. Even when it is based on sound statistical generalizations about the rate of criminal offending by members of particular racial groups, racial profiling unfairly burdens *innocent* minorities with law enforcement for the sake of promoting the greater social good (Maclin 2001; Carter 2004). For example, if the police use race as a factor when choosing which motorists to pull over for traffic stops, then *innocent* black motorists will suffer a disproportionate share of the burden of policing. Disproportionately burdening innocent members of protected groups for the sake of promoting general security is unfair to the burdened individuals, and hence unjust. Therefore, racial profiling is unjust, even when it is more effective than other policing measures at reducing crime. Call this the “Unequal Burdens Argument” against racial profiling. The argument can be spelled out in a variety of ways, but the argument’s foundational claim is that justice requires the burdens of law enforcement to be distributed roughly equally among innocent people, regardless of race, even if an unequal distribution of the burdens is optimal with respect to crime prevention.

A version of the Unequal Burdens Argument can be raised *mutatis mutandis* for predictive policing. Under a place-based predictive policing program, individuals living in designated high-risk areas receive greater attention from police. They are therefore more likely to interact with police via a stop, search, seizure, or arrest than people who live or work elsewhere. When a strong correlation exists between living in a designated high-risk area and being black, predictive policing can thus have the same outcome that people find so troubling about racial profiling: disproportionate contact with police by innocent black Americans. Increasing police contact in an area exposes innocent residents of that area to an

elevated risk of mistaken search, seizure, arrest, and conviction. These risks can lead to real harm for members of the targeted community in the form of loss of autonomy (Blitz 2004), restrictions on individual expression and creativity, violations of citizens' dignity (Castiglione 2008), and the perception of a lack of respect by members of the community (Taslitz 2003). And of course, there are well-documented risks of physical harm (Tate, Jenkins, and Rich 2020). In the extreme case, an innocent person may be killed if mistaken for an armed offender resisting arrest. There is evidence that black Americans are already 3.5 times more likely than white Americans to be the victim of a police shooting (Ross 2015). As Robertson et al (2020: 116) put the point, "Every encounter [with police] risks violating a marginalized individual's right to liberty and equality, right to be free from discrimination, and right to be free from unreasonable search and seizure."

Therefore, when predictive policing recommends a disproportionate distribution of police patrols in predominantly black communities, it disproportionately imposes these burdens on innocent members of the communities. Note that none of these burdens assumes any malicious intent or racial bias on the part of police officers. The only assumption required for the Unequal Burdens Argument to have teeth is that these burdens are more likely to eventuate when there is more interaction with police by the public. Scholars have pointed out that other burdens are less direct. Heavy policing of black communities can lead to high incarceration rates for black Americans. In 2017, black Americans constituted 12% of the U.S. population but 33% of the incarcerated population. The high proportion of black Americans with a criminal record has downstream effects on employment, which can harm community cohesion. As Tracy Meares observes, convicted

felons are less likely to invest in their own “human capital,” perceiving such investment as a waste of time given their criminal record. This frays the relations between convicts and other members of their community, because the person with a criminal record is less capable of benefiting those members (Meares and Kahan 1998). As Andrew Selbst describes the concern, “The most “accurate” decision may not lead to the fairest result for people swept up in that region” (Selbst 2017: 137)

So, even if predictive policing concentrates police patrols in minority communities simply because more crime occurs in those communities, innocent residents will be disproportionately burdened for the sake of promoting general security. This is unfair.

#### *The Unequal Benefits Objection*

The Unequal Burdens Argument against predictive policing faces a significant challenge. The challenge emerges from observing that predictive policing might disproportionately benefit the same people it burdens. Mathias Risse and Richard Zeckhauser have defended the fairness of racial profiling (in principle, if not in practice) as follows.

We submit that a sufficient condition for imposing unequal burdens...is that those burdened more are net beneficiaries from that public good [whose supply depends on their unequal burden]. If the unequal imposition of a burden is counterbalanced by a net benefit that the relevant group gains, the unequal burden is not undue (Risse and Zeckhauser 2004: 158).

David Boonin has further argued that members of black communities do in some cases enjoy a net benefit from racial profiling, at least when it is effective at reducing crime.

Boonin writes,

Black Americans, after all, are disproportionately victims of crime. When crime is reduced, therefore, black Americans on average are disproportionately the beneficiaries (Boonin 2011: 342).

Let us call this objection the *Unequal Benefits Objection* to The Unequal Burdens Argument. The objection has some merit. According to the FBI crime database, in 2015, of the 13,455 homicide cases where the FBI listed a victim's racial information, 52.3% of victims were black, despite the fact that blacks comprise just 13.3% of the U.S. population. In 2018, 70% of black victims identified the offender as black (Morgan and Oudekerk 2019). Among black victims of crime, older black Americans are most commonly the victims of crimes, and they are most commonly victimized by young black Americans. So, disproportionately targeting crime in predominantly black communities can be expected to disproportionately benefit black victims.

Even if one rejects the Unequal Benefits Objection when it is applied to racial profiling, the objection seems particularly plausible as a response to the Unequal Burdens Argument against *predictive policing*. If predictive policing is effective at reducing crime, it will be most effective at reducing crime in the places to which officers are directed. If predictive policing forecasts concentrate crime-fighting efforts in black communities, and it is effective, then predictive policing disproportionately benefits black communities in the form of crime reduction. These disproportionate benefits might morally counterbalance the

disproportionate burdens associated with additional police attention. As David Boonin writes, “If a practice disproportionately benefits black Americans, then it’s not clear why it would be unfair for it to also disproportionately burden them” (Boonin 2011: 342).

*Fairness and valid refusal of defensive intervention*

In response to the Unequal Benefits Objection I will now argue that *even if* the crime reduction benefits of predictive policing significantly outweighed the burdens for black communities, predictive policing might still be unfair. My argument begins with the observation that victims of harm possess a degree of moral authority over defensive interventions aimed at benefitting them. To see this, consider Jonathan Parry’s description of a single victim case of third-party defense against aggression.

Elevator 1: Victim is taking a ride in an elevator. Aggressor begins to saw through the elevator cable because he hates Victim and wishes him dead. Rescuer is walking by and is able to shoot Aggressor with her sniper rifle, thereby saving Victim. However, Victim refuses intervention because he is deeply committed to nonviolence (Parry 2017: 363).

It appears that it would be wrong for Rescuer to intervene in Elevator 1. Parry suggests that intervention is wrong in Elevator 1 because individuals possess a moral authority over what harms are brought against others on their behalf.

[J]ust as I can exclude you from using my body or property by withholding my consent, I have the power to prevent you from justifying your actions by appealing to the fact that you will benefit me by repudiating those benefits (Parry 2017: 371).

Parry calls this power the “power of prudential exclusion” (PPE). The moral foundation of this power is a commitment to anti-paternalism. Parry says, “It’s the victim’s life that is at stake, not anybody else’s, and so only he gets to decide if and how it is defended...On these matters, the individual is sovereign” (Parry 2017: 363). Just as it is my right to decline medical treatment because of my moral or religious commitments, it is my right to decline third-party defense that is intended to benefit me.

PPE entails that it is wrong for Rescuer to intervene in Elevator 1 because Victim’s valid refusal to consent to the intervention means that Victim’s well-being cannot be used to justify Rescuer’s harming of Aggressor. To the contrary, by invoking Victim’s well-being to justify harming Aggressor, Rescuer wrongs Victim. Cases involving *multiple* victims are morally more complex. Consider Elevator 2.

Elevator 2: Five Victims are taking a ride in an elevator. Aggressor begins to saw through the elevator cable because he hates Victims and wishes them dead. Rescuer is walking by and is able to shoot Aggressor with her sniper rifle, thereby saving Victims. However, while one Victim consents to Rescuer’s intervention, the four remaining Victims refuse because of their deep commitment to nonviolence (Parry 2017: 367–68).

PPE entails that the Rescuer cannot appeal to the benefits to the four non-consenting victims to justify killing Aggressor, but Rescuer in Elevator 2 *need not* make use of the good of the four non-consenting victims to justify killing Aggressor. The good of the one consenting victim appears *sufficient on its own* to justify Rescuer’s killing of Aggressor. Therefore, Rescuer is permitted to kill Aggressor.

Consent from potential beneficiaries of a defensive intervention is especially important when a defensive intervention imposes a risk of harm on innocent bystanders. This is because harm to innocent people is justified only if the harm prevented greatly outweighs it. As Parry puts it,

when victims validly refuse defensive intervention, they thereby exclude their good from the pool of reasons that rescuers may appeal to in order to justify defending them. However, in cases where...defense involves killing innocents, the inadmissibility of these reasons may render defensive force

disproportionate, depending on the specific facts of the case (Parry 2017: 385).

Consider an extreme case in which a defensive intervention imposes significant harms on innocent bystanders, and *all* of the beneficiaries of the intervention have validly refused the intervention. In this case, because none of the benefits to the victims count morally in favor of intervention, and the harms the intervention imposes on innocent people count strongly against it, it is morally wrong to intervene.

These observations from the ethics of third-party defense are key to understanding why the Unequal Benefits Objection to the Unequal Burdens Argument will often be unsuccessful. The practice of policing is a form of third-party defense involving multiple victims:

- (1) An aggressor (criminal) threatens unjustified harm to several victims.
- (2) A rescuer (police) has the ability to prevent the threatened harm by imposing defensive harm (arrest) on the aggressor.
- (3) Defending every victim is a condition of defending any.

And in many cases, a fourth condition will obtain as well.

(4) One subset of the members of the victim group consents to defensive intervention, and another subset refuses (Parry 2017: 360).

If we are trying to determine whether the benefits of a policing practice for members of a community are sufficient to outweigh the burdens, PPE helps us to think about which benefits *count*. PPE tells us that the benefits of a policing practice, qua defensive intervention, count morally in favor of the practice only when those benefits accrue to community members who have not validly refused consent. Disproportionate benefits accruing to beneficiaries who have refused consent cannot counteract any disproportionate burdens imposed on innocent persons. Where the subset of consenting beneficiaries is very small, and the burdens are significant, the burdens will often outweigh the benefits.

I propose that the fairness of predictive policing therefore turns in part on what proportion of innocent members of targeted communities refuse consent to predictive policing and whether this refusal is valid. I cannot settle the question about the extent to which community members have or have not refused consent to predictive policing, because no surveys exist measuring public attitudes toward predictive policing. This absence of survey evidence is troubling on its own, given what I've argued so far. There is, however, indirect evidence supporting the conclusion that a significant segment of black communities might refuse consent to predictive policing as a method of allocating police resources, especially when predictive policing allocates *more* police patrols where they live. According to a recent national Gallup poll, 81% of black Americans want the police to spend the same or more time in their communities, but 19% want the police to spend less time there (Saad

2020). Black Americans have less favorable attitudes toward police than do white Americans (Albrecht and Green 1977; Bayley and Mendelsohn 1969; Hadar and Snortum 1975; Smith 1991). Only 40% of black Americans have a favorable attitude toward police. 11% have “very unfavorable” attitudes toward police. 56% of black Americans say that police tactics are too harsh (Ekins 2016).

Suppose that a substantial number of black citizens refuse consent to predictive policing. One might question whether this refusal to consent is *valid*. If refusal to consent to defensive intervention is not valid because it is unreasonable, one’s power of prudential exclusive may not be exercised. For example, refusing to consent to a defensive intervention on the basis of trivial or morally objectionable reasons (e.g., the rescuer’s race or gender) fails to preclude one’s well-being from counting in favor of the intervention. Refusal to consent to predictive policing may appear to be unreasonable as long as it is granted that predictive policing is more effective at apprehending and deterring offenders than alternative methods. Citizens of legitimate states have civic obligations to respect the authority of the law, where an obligation to respect the authority of the law includes an obligation to acquiesce to effective measures of crime reduction. If such an obligation exists, then it may seem doubtful that a person, qua citizen of a legitimate state, possesses the normative power to prohibit a state agency from pursuing the democratically authorized aim of crime reduction.

To the contrary, black Americans can have powerful prudential and moral reasons to refuse predictive policing, and these reasons might override their civic obligation to acquiesce to effective state-authorized crime prevention initiatives. I will now describe some of these reasons.

First, many black Americans—particularly individuals who have been confined to ghetto poverty—may have diminished civic obligations to uphold the law as a general matter. Civic obligations are those which arise from the bonds between citizens in virtue of which they together compose a polity or nation. Were these associational ties between citizens not to exist, nor would civic obligations. Thomas Shelby has argued that black Americans living in ghetto poverty have diminished civic obligations to respect the authority of the law insofar as civic obligations are rooted in the political value of reciprocity (Shelby 2007: 145). In virtue of benefitting from the cooperative scheme embodied by our social institutions, citizens incur obligations not to exploit or take advantage of other citizens in ways that flout the demands of that cooperative scheme. But, as Shelby observes, if reciprocity is the grounds of civic obligation, then it is plausible that the strength of one's civic obligation depends on society's institutions being arranged in such a way that one can expect to enjoy their fair share of benefits from the cooperative scheme. And all empirical facts about the past and present state of black America with respect to economic well-being, education, employment, crime and criminality, and life-expectancy indicate that American social institutions continue to be arranged in such a way that black Americans are deprived their fair share of benefits from the scheme of cooperation. Shelby writes, "From the standpoint of many ghetto residents, the social order lacks legitimacy. There appears to be a conspiracy to contain, exploit, and underdevelop the black urban poor, to deny them equal civic standing and punish them when they refuse to accommodate themselves to injustice" (Shelby 2007: 170). If one's civic obligations are dependent on the reasonable expectation that the basic structure is arranged so that one will receive their fair share of benefits of the

system of cooperation, black Americans living in ghetto poverty possess weakened obligations to uphold the law.

Second, black Americans have significant reasons to distrust federal criminal justice policies. Members of black communities have anecdotal, historical, media-based, and empirical evidence that the concentration of police patrols in their communities will impose significant moral costs on them and people with whom they are intimately related. Studies show racial disparities in police use of force (Buehler 2017), drug enforcement (Lynch et al. 2013), and traffic ticketing (Dunn 2009). 39% of black Americans know someone who has been physically mistreated by the police, which is twice the percentage of white Americans. Black Americans comprised just 13 percent of the U.S. population but 25% of all police killings and 1/3 of unarmed victims (Fox et al. 2019). Moreover, given the role that law enforcement played in upholding institutions of which black Americans were the primary victims, including slavery, segregation, and Jim Crow, it is reasonable for black Americans to see these statistical disparities as reflecting a continued discounting of their interests by criminal justice institutions. So, when presented with a novel policing practice such as predictive policing, one result of which is the concentration of police patrols in black communities, it is prudentially reasonable for black Americans to refuse consent on the grounds that the practice is likely to be to their detriment. The combination of weakened civic obligations and a background of legitimate distrust of criminal justice institutions establishes a strong *prima facie* case for the claim that it is reasonable for black Americans to refuse consent to predictive policing.

Finally, one particularly powerful reason to refuse defensive intervention is that the intervention threatens harm to an innocent person to whom the beneficiary of the intervention owes an associative or role-based obligation. Members of black communities can find themselves in precisely this situation with respect to police protection.

An associative obligation is a moral obligation that a person has in virtue of standing in a special relation to someone else. On associative obligations, C.D. Broad writes,

[E]ach of us has specially urgent obligations to benefit certain individuals and groups which stand in certain special relations to himself, e.g., his parents, his children, his fellow-countrymen, etc. [T]hese special relationships are the ultimate and sufficient ground for these specially urgent claims on one's beneficence (Broad 1985: 220).

Ordinary morality includes a commitment to the view that individuals possess especially strong *associative obligations* to individuals with whom they share a special relationship—family members in particular—and that we ought to give greater weight to the interests of friends and family in our moral deliberations than we give to the interests of strangers. *Role obligations* are those obligations that one occupies in virtue of occupying a particular social role. For instance, a firefighter has an obligation to respond to an emergency in virtue of their occupying the role of firefighter. In ordinary circumstances, a parent plausibly has both associative and role obligations to protect their children from harm. About familial obligations Michael Hardimon writes,

What, after all, would it mean to say that we conceive of ourselves as family members but regard ourselves as having no familial obligations whatsoever?

The suggestion that we could do this is incoherent...Although...

noncontractual role obligations are by no means the whole of family or civic life, they are an integral part of them (Hardimon 1994: 346).

If an action, practice, or policy were to conflict with a person's associative or role obligations to family members, this would constitute one powerful reason to reject it.

Many Black Americans have reasons, grounded in associative and role duties to friends and family to refuse policing practices that concentrate police patrols in their communities. As an illustration, consider once again the fact that among Black Americans, young men are by far the most likely to be the victims of police violence (Fox et al. 2019). Black Americans who are parents and grandparents have associative and role obligations to their children and grandchildren to protect them from harm. And so, many Black Americans who are parents and grandparents can reasonably view the intensification of police attention in their communities as being in tension with their associative and role obligations, *qua* parents and grandparents, to protect their children and grandchildren from harm. And if it is at least *pro tanto* reasonable for them to see a tension between intensified police patrols in their community and their familial obligations, then it is at least *pro tanto* reasonable for them to withhold consent to any practice or policy that would concentrate police patrols in their community. Predictive policing is one such practice.

In sum, I have argued that a combination of factors makes it reasonable for black Americans to refuse consent to predictive policing. Black Americans (a) possess weakened civic obligations to uphold the law, (b) are reasonably distrustful of American criminal justice practices, and (c) possess associative and role obligations to protect their loved ones from

threats of harm by police. Therefore, responding to the Unequal Burdens Argument against predictive policing requires more than showing that predictive policing produces unequal benefits in crime reduction for the same communities that it unequally burdens. It must also be shown that enough of the beneficiaries do not validly refuse consent to predictive policing. Police attention directed at a community must be *welcome* to be fair.

#### **4 Community-led policing and problem-oriented policing**

I want to close by considering two modest steps that police departments might take to address the concerns I have raised. Research shows that perceptions of police become more favorable when police focus on serious crimes rather than nuisance crimes (Maxson, Hennigan, and Sloane 2003). If a convincing case is made to members of black communities that predictive policing targets serious offenses (property crime and violent crime) rather than minor offenses, and that it can reduce crime *without* imposing undue costs on innocent people, this is one step toward securing consent. Evidence also suggests that favorable attitudes toward police are positively correlated with informal contact with police (e.g., conversing in public or in a meeting of community members) and negatively correlated with formal contact (e.g. traffic stop or arrest) (Maxson, Hennigan, and Sloane 2003). What can be done to (a) make the case to black communities that predictive policing targets serious crimes and (b) increase informal (but not formal) contact between police and community members? Understanding that any policy proposal faces empirical questions that this author is ill-equipped to answer, I very tentatively endorse a two-pronged strategy of community-led, problem-oriented policing.

*Community-led policing*

Community-led policing includes members of the community in strategic decision-making. This requires police to meet with community members to exchange feedback about priorities, strategies, and outcomes. When communities are kept abreast of, and even involved in, strategic decision-making, buy-in is a more attainable goal.

One obstacle to community support is that most predictive policing algorithms are inscrutable or “opaque” to both the police who employ them in their work and to the citizens they affect. There are three ways that predictive algorithms might be opaque: the code may not be made publicly available because of intellectual private property protections and concerns about competitive advantage; understanding programming languages is a specialized skill, and few non-programmers are computationally literate in ways that would allow them to understand why an algorithm makes the determinations that it does; and the decision procedures of machine learning algorithms, which work by a mathematical process of iterative statistical optimization, resist interpretation in terms comprehensible to a human (Burrell 2016). Algorithmic opacity threatens community buy-in through a combination of factors. First, it can call into question the trustworthiness of the agency deploying the algorithm by presenting *opportunities for concealment* of abuses: (1) opacity enables “discrimination masking,” wherein the true (possibly racist) motives of law enforcement agents are hidden behind a mathematical façade (Barocas and Selbst 2016); (2) opacity can give rise to the appearance of “responsibility gaps,” scenarios where no one is accountable for mistakes made by the algorithmic system (Matthias 2004); (3) opacity allows for the possibility of so-called “agency laundering” wherein agents attempt to defer responsibility

for morally significant decisions to the determinations of algorithmic systems (Rubel, Castro, and Pham 2019).

At the same time, opacity makes it difficult for police to address these grounds of community distrust because police are unable to *explain* the method by which the predictive system produces crime forecasts. Opacity therefore precludes departments from addressing some of the grounds for community distrust, by preventing a free and open exchange. For these reasons algorithmic opacity undermines attempts to secure community endorsement of predictive policing.

On the other hand, the decision-making of human crime analysts is perhaps less transparent than the procedure employed by a predictive policing system. Human analysts are influenced by their own biases when making judgments about crime risk, and there is no way to assess how these biases influence those judgments. In this respect, algorithms are more transparent than human decision-makers, at least in cases where we have information about the system's input data. Evidence-based policing expert Jerry Ratcliffe writes, "Unlike with clinical or hybrid decision-making, at least administrators and voters can make an informed decision regarding the algorithms and data that are inputs for predictive policing algorithms" (Ratcliffe 2019: 360). And yet there remains an important sense in which the decision-making of human crime analysts is more accountable than decision-making assisted by algorithmic predictive policing systems: when risk assessments are made by a human analyst, one can sensibly demand a *justification* from the analyst for the methods employed. Indeed, Rachel Boba Santos suggests that this may explain why organizations like the ACLU have singled out predictive policing for criticism but not the entire enterprise of crime

analysis. She writes, "...predictive policing is not transparent and does not include a "human factor" (i.e. crime analyst) through which analytical decisions are made about data, and the appropriateness of analytical techniques [is] justified" (Boba Santos 2019: 391).

### *Problem-oriented policing*

Especially when consent is hard to secure, problem-oriented policing can secure buy-in while reducing the burdens that policing imposes on a community. Problem-oriented policing involves identifying the underlying causes of crime incidents at a hot spot and then developing and testing interventions that address those causes, rather than focusing solely on the incidents themselves (A. A. Braga 2008). Problem-oriented policing can reduce crime by eliminating opportunities for crime without concentrating police patrols in a community, and it has been found to have a statistically significant impact on crime and disorder (Weisburd et al. 2010). Because it considers non-enforcement-oriented interventions aimed at addressing the causes of crime hot spots, which do not involve increasing contact between citizens and officers on patrol, problem-oriented policing mitigates the risk of harm to innocents. Because it lowers the risk to innocents, problem-oriented policing is easier to justify in the face of concerns about fairness than enforcement-oriented policing tactics, especially in scenarios where the consent of citizens cannot be taken for granted. In addition to addressing the concern about unfairness, problem-oriented measures are also supported by a simple harm-minimization principle. According to such a principle, one is obligated to pursue the least harmful course of action in pursuing some legitimate policing aim. Because they mitigate harm to innocent citizens, non-enforcement, problem-oriented measures should be preferred to intensified patrols where they would achieve similar levels of

deterrence. This is *especially* true when community relations with police are fraught. Police departments should therefore adopt predictive policing systems that are amenable to problem-oriented policing tactics.

I have argued that the prevailing focus on racial bias in predictive policing systems has overshadowed two normative factors that are essential to a full assessment of the moral permissibility of predictive policing: fairness in the social distribution of the benefits and burdens of policing as well as the distinctive role of consent in determining fair distributions. I have proposed that individuals possess a normative power to preclude benefits to themselves from being used to justify harms imposed on others. This means that the benefits of predictive policing conferred on members of communities who have validly refused consent to predictive policing cannot be used to morally justify the extra burdens that predictive policing imposes on other innocent members of those communities. In some communities of color, where relations with police are fraught and consent is hard to secure, the threat of unfairness looms large. Achieving fairness in predictive policing will therefore often require either securing a greater degree of consent from affected communities or reducing the burdens that predictive policing imposes on those communities. For these reasons, community-led policing and problem-oriented policing are likely to be key components of any fair approach to predictive policing.

## References

- Albrecht, Stan, and Miles Green. (1977) 'Attitudes Toward the Police and the Larger Attitude Complex.' *Criminology*, 15 (1), 67–86. <https://doi.org/10.1111/j.1745-9125.1977.tb00049.x>.
- Anderson, Elijah. (2000) *Code of the Street: Decency, Violence, and the Moral Life of the Inner City*. Reprint edition. New York, NY: W. W. Norton & Company.
- Barocas, Solon, and Andrew D. Selbst. (2016) 'Big Data's Disparate Impact.' *California Law Review*, 104 (3), 671–732.
- Bayley, David H., and Harold Mendelsohn. (1969) *Minorities and the Police: Confrontation in America*. First edition. Free Press.
- Black, Donald J. (1970) 'Production of Crime Rates.' *American Sociological Review*, 35 (4), 733–48. <https://doi.org/10.2307/2093948>.
- Blitz, Marc J. (2004) 'Video Surveillance and the Constitution of Public Space: Fitting the Fourth Amendment to a World That Tracks Image and Identity.' *Texas Law Review*, 82, 1349–1481.
- Boba Santos, Rachel. (2019) 'Predictive Policing: Where's the Evidence?' in *Police Innovation: Contrasting Perspectives*, eds. David Weisburd and Anthony A. Braga, 2nd edition, 366–98. Cambridge, United Kingdom; New York, NY: Cambridge University Press.
- Boonin, David. (2011) *Should Race Matter?: Unusual Answers to the Usual Questions*. Cambridge: Cambridge University Press. <https://doi.org/10.1017/CBO9781139003650>.
- Braga, Anthony A. (2008) *Problem-Oriented Policing and Crime Prevention, 2nd Edition*. 2nd edition. New York: Lynne Rienner Publishers.
- Braga, Anthony A., Andrew V. Papachristos, and David M. Hureau. (2014) 'The Effects of Hot Spots Policing on Crime: An Updated Systematic Review and Meta-Analysis.' *Justice Quarterly*, 31 (4), 633–63. <https://doi.org/10.1080/07418825.2012.673632>.
- Braga, Anthony, Andrew Papachristos, and David Hureau. (2012) 'Hot Spots Policing Effects on Crime.' *Campbell Systematic Reviews*, 8 (1), 1–96. <https://doi.org/10.4073/csr.2012.8>.
- Brantingham, P. Jeffrey, Matthew Valasik, and George O. Mohler. (2018) 'Does Predictive Policing Lead to Biased Arrests? Results From a Randomized Controlled Trial.' *Statistics and Public Policy*, 5 (1), 1–6. <https://doi.org/10.1080/2330443X.2018.1438940>.
- Broad, C. D. (1985) *Ethics*. Edited by C. Lewy. Nijhoff International Philosophy Series. Springer Netherlands. <https://doi.org/10.1007/978-94-009-5057-3>.
- Buehler, James W. (2017) 'Racial/Ethnic Disparities in the Use of Lethal Force by US Police, 2010–2014.' *American Journal of Public Health*, 107 (2), 295–97. <https://doi.org/10.2105/AJPH.2016.303575>.
- Burrell, Jenna. (2016) 'How the Machine 'Thinks': Understanding Opacity in Machine Learning Algorithms.' *Big Data & Society*, 3 (1), 1-12. <https://doi.org/10.1177/2053951715622512>.
- Caplan, Joel, Leslie Kennedy, and Grant Drawve. (2017) 'Risk-Based Policing in Atlantic City 2016 Report.' [http://www.rutgerscps.org/uploads/2/7/3/7/27370595/acpd\\_rbpilot\\_2016reportbrief.pdf](http://www.rutgerscps.org/uploads/2/7/3/7/27370595/acpd_rbpilot_2016reportbrief.pdf).
- Carter, William. (2004) 'A Thirteenth Amendment Framework for Combating Racial Profiling.' *Harvard Civil Liberties Law Review*, 39 (1), 17–94.
- Castiglione, John D. (2008) 'Human Dignity under the Fourth Amendment.' *Wisconsin Law Review*, 4, 655–712.

- Dunn, Ronnie A. (2009) 'Measuring Racial Disparities in Traffic Ticketing within Large Urban Jurisdictions.' *Public Performance & Management Review*, 32 (4), 537–61.
- Ekins, Emily. (2016) 'Policing in America: Understanding Public Attitudes Toward the Police. Results from a National Survey.' *SSRN Electronic Journal*, January. <https://doi.org/10.2139/ssrn.2919449>.
- Ensign, Danielle, Sorelle A. Friedler, Scott Neville, Carlos Scheidegger, and Suresh Venkatasubramanian. (2018) 'Runaway Feedback Loops in Predictive Policing.' In *Conference on Fairness, Accountability and Transparency*, 81,160–71. PMLR. <http://proceedings.mlr.press/v81/ensign18a.html>.
- Ferguson, Andrew Guthrie. (2017) *The Rise of Big Data Policing: Surveillance, Race, and the Future of Law Enforcement*. New York: NYU Press.
- Fox, Joe, Adrian Blanco, Jennifer Jenkins, Julie Tate, and Wesley Lowery. (2019) 'What We've Learned about Police Shootings 5 Years after Ferguson.' *Washington Post*, August 9, 2019, sec. National. <https://www.washingtonpost.com/nation/2019/08/09/what-weve-learned-about-police-shootings-years-after-ferguson/>.
- Gutierrez, Carmen M., and David S. Kirk. (2017) 'Silence Speaks: The Relationship between Immigration and the Underreporting of Crime.' *Crime & Delinquency*, 63 (8), 926–50. <https://doi.org/10.1177/0011128715599993>.
- Hadar, Ilana, and John R. Snortum. (1975) 'The Eye of the Beholder: Differential Perceptions of Police by the Police and the Public.' *Correctional Psychologist*, 2 (1), 37–54. <https://doi.org/10.1177/009385487500200103>.
- Harcourt, Bernard E. (2006) *Against Prediction: Profiling, Policing, and Punishing in an Actuarial Age*. Chicago: University of Chicago Press.
- Hardimon, Michael O. (1994) 'Role Obligations.' *The Journal of Philosophy*, 91 (7), 333–63. <https://doi.org/10.2307/2940934>.
- Haskins, Caroline. (2019) 'Academics Confirm Major Predictive Policing Algorithm Is Fundamentally Flawed.' *Vice*. February 14, 2019. <https://www.vice.com/en/article/xwbag4/academics-confirm-major-predictive-policing-algorithm-is-fundamentally-flawed>.
- "ILP Manual." (2016) Pasco Sheriff's Office. [https://www.flsheriffs.org/uploads/docs/ILP\\_manual\\_final\\_edit\\_010716.pdf](https://www.flsheriffs.org/uploads/docs/ILP_manual_final_edit_010716.pdf).
- Jouvenal, Justin. (2016) 'Police Are Using Software to Predict Crime. Is It a 'Holy Grail' or Biased against Minorities?' *Washington Post*, November 17, 2016, sec. Public Safety. [https://www.washingtonpost.com/local/public-safety/police-are-using-software-to-predict-crime-is-it-a-holy-grail-or-biased-against-minorities/2016/11/17/525a6649-0472-440a-aae1-b283aa8e5de8\\_story.html](https://www.washingtonpost.com/local/public-safety/police-are-using-software-to-predict-crime-is-it-a-holy-grail-or-biased-against-minorities/2016/11/17/525a6649-0472-440a-aae1-b283aa8e5de8_story.html).
- Kirk, David S., and Andrew V. Papachristos. (2011) 'Cultural Mechanisms and the Persistence of Neighborhood Violence.' *American Journal of Sociology*, 116 (4), 1190–1233. <https://doi.org/10.1086/655754>.
- Kump, Paul, David Haro Alonso, Yongyi Yang, Joseph Candella, Jonathan Lewin, and Miles N. Wernick. (2016) 'Measurement of Repeat Effects in Chicago's Criminal Social Network.' *Applied Computing and Informatics*, 12 (2), 154–60. <https://doi.org/10.1016/j.aci.2016.01.002>.
- Lau, Tim. (2020) 'Predictive Policing Explained | Brennan Center for Justice.' April 1, 2020. <https://www.brennancenter.org/our-work/research-reports/predictive-policing-explained>.

- Levine, E. S., Jessica Tisch, Anthony Tasso, and Michael Joy. (2017) 'The New York City Police Department's Domain Awareness System.' *INFORMS Journal on Applied Analytics*, 47 (1), 70–84. <https://doi.org/10.1287/inte.2016.0860>.
- Lum, Kristian, and William Isaac. (2016) 'To Predict and Serve?' *Significance*, 13 (5), 14–19. <https://doi.org/10.1111/j.1740-9713.2016.00960.x>.
- Lynch, Mona, Marisa Omori, Aaron Roussell, and Matthew Valasik. (2013) 'Policing the 'Progressive' City: The Racialized Geography of Drug Law Enforcement.' *Theoretical Criminology*, 17 (3), 335–57. <https://doi.org/10.1177/1362480613476986>.
- Maclin, Tracey. (2001) 'The Fourth Amendment on the Freeway Symposium: Racial Profiling: A New Road Hazard.' *Rutgers Race & the Law Review*, 3 (2), 117–90.
- Matthias, Andreas. (2004) 'The Responsibility Gap: Ascribing Responsibility for the Actions of Learning Automata.' *Ethics and Information Technology*, 6 (3), 175–83. <https://doi.org/10.1007/s10676-004-3422-1>.
- Maxson, Cheryl, Karen Hennigan, and David Sloane. (2003) 'Factors That Influence Public Opinion of the Police.' 197925. National Institute of Justice. <https://nij.ojp.gov/library/publications/factors-influence-public-opinion-police>.
- Meares, Tracey L., and Dan M. Kahan. (1998) 'Law and (Norms of) Order in the Inner City.' *Law & Society Review*, 32 (4), 805–38. <https://doi.org/10.2307/827740>.
- Meijer, Albert, and Martijn Wessels. (2019) 'Predictive Policing: Review of Benefits and Drawbacks.' *International Journal of Public Administration*, 42 (12), 1031–39. <https://doi.org/10.1080/01900692.2019.1575664>.
- Miller, Leila. (2020) 'LAPD Will End Controversial Program That Aimed to Predict Where Crimes Would Occur.' *Los Angeles Times*, April 21, 2020, sec. California. <https://www.latimes.com/california/story/2020-04-21/lapd-ends-predictive-policing-program>.
- Moravec, Eva Ruth. (2019) 'Do Algorithms Have a Place in Policing?' *The Atlantic*, September 5, 2019. <https://www.theatlantic.com/politics/archive/2019/09/do-algorithms-have-place-policing/596851/>.
- Morgan, Rachel, and Barbara Oudekerk. (2019) 'Bureau of Justice Statistics (BJS) - Criminal Victimization, 2018.' 253043. Criminal Victimization Series. Bureau of Justice. <https://www.bjs.gov/index.cfm?ty=pbdetail&iid=6686>.
- National Research Council. (2004) *Fairness and Effectiveness in Policing: The Evidence*. Washington, DC: The National Academies Press. <https://doi.org/10.17226/10419>.
- O'Neil, Cathy. (2016) *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. 1st edition. New York: Crown.
- Parry, Jonathan. (2017) 'Defensive Harm, Consent, and Intervention.' *Philosophy & Public Affairs*, 45 (4), 356–96. <https://doi.org/10.1111/papa.12099>.
- Perry, Walter L., Brian McInnis, Carter C. Price, Susan Smith, and John S. Hollywood. (2013) 'Predictive Policing: The Role of Crime Forecasting in Law Enforcement Operations.' RR-233-NIJ. RAND Corporation. [https://www.rand.org/pubs/research\\_reports/RR233.html](https://www.rand.org/pubs/research_reports/RR233.html).
- Pierce, Glenn L., Susan Spaar, and LeBaron R Briggs. (1988) *The Character of Police Work: Strategic and Tactical Implications*. Boston, Mass.: Center for Applied Social Research, Northeastern University.

- Ratcliffe, Jerry. (2019) 'Advocate: Predictive Policing.' in *Police Innovation: Contrasting Perspectives*, eds. David Weisburd and Anthony A. Braga, 2nd edition, 347–65. Cambridge, United Kingdom; New York, NY: Cambridge University Press.
- Richardson, Rashida, Jason Schultz, and Kate Crawford. (2019) 'Dirty Data, Bad Predictions: How Civil Rights Violations Impact Police Data, Predictive Policing Systems, and Justice.' *NYU Law Review*, 94 (192), 15–55.
- Risse, Mathias, and Richard Zeckhauser. (2004) 'Racial Profiling.' *Philosophy & Public Affairs*, 32 (2), 131–70. <https://doi.org/10.1111/j.1088-4963.2004.00009.x>.
- Ross, Cody T. (2015) 'A Multi-Level Bayesian Analysis of Racial Bias in Police Shootings at the County-Level in the United States, 2011–2014.' *PLoS ONE*, 10 (11), e0141854. <https://doi.org/10.1371/journal.pone.0141854>.
- Rubel, Alan, Clinton Castro, and Adam Pham. (2019) 'Agency Laundering and Information Technologies.' *Ethical Theory and Moral Practice*, 22 (4), 1017–41. <https://doi.org/10.1007/s10677-019-10030-w>.
- Saad, Lydia. (2020) 'Black Americans Want Police to Retain Local Presence.' Gallup.com. August 5, 2020. <https://news.gallup.com/poll/316571/black-americans-police-retain-local-presence.aspx>.
- Selbst, Andrew D. (2017) 'Disparate Impact in Big Data Policing.' *Georgia Law Review*, 52 (1), 109–195.
- Shapiro, Aaron. (2017) 'Reform Predictive Policing.' *Nature*, 541 (7638), 458–460. <https://doi.org/10.1038/541458a>.
- Shelby, Tommie. (2007) 'Justice, Deviance, and the Dark Ghetto.' *Philosophy & Public Affairs*, 35 (2), 126–60.
- Sherman, Lawrence W., Patrick R. Gartin, and Michael E. Buerger. (1989) 'Hot Spots of Predatory Crime: Routine Activities and the Criminology of Place\*.' *Criminology*, 27 (1), 27–56. <https://doi.org/10.1111/j.1745-9125.1989.tb00862.x>.
- Smith, David J. (1991) 'The Origins of Black Hostility to the Police.' *Policing and Society*, 2 (1), 1–15. <https://doi.org/10.1080/10439463.1991.9964628>.
- Tashea, Jason. (2016) 'Websites and Apps for Sharing Crime and Safety Data Have Become Outlets for Racial Profiling.' *ABA Journal*, August 1, 2016. [https://www.abajournal.com/magazine/article/crime\\_safety\\_website\\_racial\\_profiling](https://www.abajournal.com/magazine/article/crime_safety_website_racial_profiling).
- Taslitz, Andrew. (2003) 'Respect and the Fourth Amendment.' *Journal of Criminal Law and Criminology*, 94 (1), 15–102.
- Tate, Julie, Jennifer Jenkins, and Stephen Rich. (2020) 'Fatal Force: 2019 Police Shootings Database.' Washington Post. August 10, 2020. <https://www.washingtonpost.com/graphics/2019/national/police-shootings-2019/>.
- Vitale, Alex S. (2017) *The End of Policing*. London ; New York: Verso.
- Weisburd, David, and John E. Eck. (2004) 'What Can Police Do to Reduce Crime, Disorder, and Fear?' *The ANNALS of the American Academy of Political and Social Science*, 593 (1), 42–65. <https://doi.org/10.1177/0002716203262548>.
- Weisburd, David, Lisa Maher, Lawrence Sherman, Michael Buerger, Ellen Cohn, and Anthony Petrosino. (1993) 'Contrasting Crime General and Crime Specific Theory: The Case of Hot Spots of Crime.' in *Advances in Criminological Theory*, 4, 45–70. Routledge.
- Weisburd, David, Cody W. Telep, Joshua C. Hinkle, and John E. Eck. (2010) 'Is Problem-Oriented Policing Effective in Reducing Crime and Disorder?' *Criminology & Public Policy*, 9 (1), 139–72. <https://doi.org/10.1111/j.1745-9133.2010.00617.x>.